

TABLE OF CONTENTS

I. REAL PARTY IN INTEREST	1
II. RELATED APPEALS AND INTERFERENCES	1
III. STATUS OF CLAIMS	2
IV. STATUS OF AMENDMENTS	2
V. SUMMARY OF CLAIMED SUBJECT MATTER	2
VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL	3
VII. ARGUMENT	4
VIII. CLAIMS APPENDIX.....	17
IX. EVIDENCE APPENDIX.....	21
X. RELATED PROCEEDINGS APPENDIX.....	22

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES**

In re Application of	:	Customer Number: 46320
	:	
David KAMINSKY, et al.	:	Confirmation Number: 3989
	:	
Application No.: 10/663,125	:	Group Art Unit: 2152
	:	
Filed: September 16, 2003	:	Examiner: P. Lee
	:	
For: AUTONOMIC CLUSTER-BASED OPTIMIZATION	:	

APPEAL BRIEF

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This Appeal Brief is submitted in support of the Notice of Appeal filed September 11, 2008, wherein Appellants appeal from the Examiner's rejection of claims 6-17.

I. REAL PARTY IN INTEREST

This application is assigned to IBM Corporation by assignment recorded on September 16, 2003, at Reel 014516, Frame 0548.

II. RELATED APPEALS AND INTERFERENCES

Appellants are unaware of any related appeals and interferences.

III. STATUS OF CLAIMS

Claims 1-17 are pending, and claims 6-17 are two-times rejected in this Application. Claims 1-5 have been withdrawn from consideration pursuant to the provisions of 37 C.F.R. § 1.142(b). It is from the multiple rejections of claims 6-17 that this Appeal is taken.

IV. STATUS OF AMENDMENTS

The claims have not been amended subsequent to the imposition of the Second and Final Office Action dated June 11, 2008 (hereinafter the Second Office Action).

V. SUMMARY OF CLAIMED SUBJECT MATTER

Referring to Figure 2 and also to independent claim 6, a method for autonomically optimizing a cluster of nodes is disclosed. In block 205, a node in the cluster which requires re-configuration is detected (page 13, lines 10-16). In block 215, a workload hosted by the node is identified, and in block 220, a set of configuration parameters associated with the workload is retrieved (page 13, line 19 through page 14, line 3). In block 225, a new generation of configuration parameters is produced based upon the retrieved set using a genetic computing process (page 14, lines 4-5). In block 255, the node is reconfigured with selected ones of the new generation of configuration parameters (page 14, lines 17-19).

Referring to Figure 2 and also to independent claim 12, a machine readable storage having stored thereon a computer program for autonomically optimizing a cluster of nodes is disclosed. The computer program comprises a routine set of instructions for causing the machine to perform the following steps. In block 205, a node in the cluster which requires re-configuration is detected (page 13, lines 10-16). In block 215, a workload hosted by the node is

1 identified, and in block 220, a set of configuration parameters associated with the workload is
2 retrieved (page 13, line 19 through page 14, line 3). In block 225, a new generation of
3 configuration parameters is produced based upon the retrieved set using a genetic computing
4 process (page 14, lines 4-5). In block 255, the node is reconfigured with selected ones of the
5 new generation of configuration parameters (page 14, lines 17-19).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

1. Claims 7 and 13 were rejected under the second paragraph of 35 U.S.C. § 112; and
2. Claims 6-17 were rejected under 35 U.S.C. § 103 for obviousness based upon Maltz et al., U.S. Patent Publication No. 2002/0143929 (hereinafter Maltz), in view of Nozawa et al., U.S. Patent No. 6,272,543 (hereinafter Nozawa).

VII. ARGUMENT

THE REJECTION OF CLAIMS 7 AND 13 UNDER THE SECOND PARAGRAPH OF 35 U.S.C. §

112

For convenience of the Honorable Board in addressing the rejections, claim 13 stands or falls together with independent claim 7.

Appellants respectfully submit that the Examiner has failed to establish a prima facie case of indefiniteness under the second paragraph of 35 U.S.C. § 112. M.P.E.P. § 2173.02 states the following:

If upon review of a claim in its entirety, the examiner concludes that a rejection under 35 U.S.C. 112, second paragraph, is appropriate, such a rejection should be made and an analysis as to why the phrase(s) used in the claim is "vague and indefinite" should be included in the Office action. (emphasis added).

As stated in Metabolite Labs., Inc. v. Lab. Corp. of Am. Holdings,¹ "[o]nly when a claim remains insolubly ambiguous without a discernible meaning after all reasonable attempts at construction must a court declare it indefinite." The Examiner, however, has not forth an analysis as to why the phrase at issue is vague and indefinite.

As noted by Appellants on page 8, lines 5-7 of the First Response dated February 20, 2008 (hereinafter the First Response), the recitation of "the group" in claims 7 and 13, which is the subject of the Examiner's rejection, is a commonly used phrase in claims. In this regard, reference is made to M.P.E.P. § 2173.05(h) as to Markush Claims. As evident from this first sentence, the format "selected from the group consisting of ..." has consistently been considered

¹ 370 F.3d 1354, 1366, 71 USPQ2d 1081, 1089 (Fed. Cir. 2004).

as proper claim language by the Patent Office. Appellants' position, therefore, is that one having ordinary skill in the art would have no difficulty understanding the scope of claims 7 and 13.

Since the Examiner has failed to establish a prima facie case of indefiniteness and since the language at issue has long been accepted by the Patent Office as being proper, Appellants respectfully submit that the Examiner has committed error in rejecting claims 7 and 13 under the second paragraph of 35 U.S.C. § 112.

THE REJECTION OF CLAIMS 6-17 UNDER 35 U.S.C. § 103 FOR OBVIOUSNESS BASED UPON MALTZ IN VIEW OF NOZAWA

For convenience of the Honorable Board in addressing the rejections, claims 7-17 stand or fall together with independent claim 6.

As is evident from Appellants' previously-presented comments during prosecution of the present Application and from Appellants' comments below, there are questions as to how the limitations in the claims correspond to features in the applied prior art. In this regard, reference is made to M.P.E.P. § 1207.02, entitled "Contents of Examiner's Answer." Specifically, the following is stated:

(A) CONTENT REQUIREMENTS FOR EXAMINER'S ANSWER. The examiner's answer is required to include, under appropriate headings, in the order indicated, the following items:

...

(9)(c) For each rejection under 35 U.S.C. 102 or 103 where there are questions as to how limitations in the claims correspond to features in the prior art even after the examiner complies with the requirements of paragraphs (c) and (d) of this section, the examiner must compare at least one of the rejected claims feature by feature with the prior art relied on in the rejection. The comparison must align the language of the claim side-by-side with a reference to the specific page, line number, drawing reference number, and quotation from the prior art, as appropriate. (emphasis added)

Therefore, if the Examiner is to maintain the present rejections and intends to file an Examiner's Answer, the Examiner is required to include the aforementioned section in the Examiner's Answer.

Appellants have compared the statement of the rejection found on pages 3-6 of the Second Office Action with the statement of the rejection found on pages 4-7 of the First Office Action. Upon making this comparison, Appellants have been unable to discover any substantial differences between the respective statements of the rejection. As such, Appellants proceed on the basis that the Examiner's sole response to the arguments presented in Appellants' First Response is found on pages 8-10 of the Second Office Action in the section entitled "Response to Arguments."

Obviousness is a legal conclusion based on underlying facts of four general types, all of which must be considered by the trier of fact: (1) the scope and content of the prior art; (2) the level of ordinary skill in the art; (3) the differences between the claimed invention and the prior art; and (4) any objective indicia of nonobviousness.² Upon reviewing the Examiner's statement of the rejection with regard to claims 6 and 12 on page 5 of the First Office Action, Appellants submit that the Examiner has failed to clearly designate the teachings in Maltz being relied upon the statement of the rejection. In this regard, the Examiner's rejection under 35 U.S.C. § 103 also fails to comply with 37 C.F.R. § 1.104(c), which provides:

In rejecting claims for want of novelty or for obviousness, the examiner must cite the best references at his or her command. When a reference is complex or shows or describes inventions other than that claimed by the applicant, the particular part relied on must be designated as nearly

² See *KSR Int'l v. Teleflex Inc.*, 550 U.S. ____ (2007); *Graham v. John Deere Co.*, 383 U.S. 1, 17-18 (1966); *Continental Can Co. USA, Inc. v. Monsanto Co.*, 948 F.2d 1264, 1270, 20 USPQ2d 1746, 1750-51 (Fed. Cir. 1991); *Panduit Corp. v. Dennison Mfg. Co.*, 810 F.2d 1561, 1566-68, 1 USPQ2d 1593, 1594 (Fed. Cir. 1987).

as practicable. The pertinence of each reference, if not apparent, must be clearly explained and each rejected claim specified.

The Examiner has not designated the particular parts of Maltz being relied "as nearly as practicable." Instead, the Examiner's statement of rejection reproduces entire passages of the claims and asserts that all the limitations in the particular passages are disclosed by certain cited passages within Maltz. Although broadly identifying where the Examiner believes the claimed limitations are disclosed in Maltz, the statement of the rejection fails to clearly identify many of the specific elements within Maltz being relied upon in the rejection, as required by 37 C.F.R. § 1.104(c).

By not clearly indicating those specific elements being relied upon in the prior art to teach the invention, as recited in the claims, the Examiner has failed to fully establish the underlying facts regarding (1) the scope and content of the prior art and (3) the differences between the claimed invention and the prior art. Thus, the Examiner has improperly arrived at the legal conclusion that the claimed invention is obvious based upon the combination of Maltz in view of Nozawa.

The above-reproduced arguments (incorporated herein) were originally presented on page 9, line 1 through page 10, line 3 of the First Response. However, the Examiner's did not directly respond to these arguments in the Second Office Action.

Notwithstanding the Examiner's failure to clearly designate the teachings in Maltz being relied upon in the statement of the rejection, the Examiner has failed to properly establishing the

underlying facts regarding the scope and content of the prior art and the differences between the claimed invention and the prior art. For example, regarding the claimed "detecting a node in the cluster which requires re-configuration," the Examiner asserted "detecting scheduled collection/transmission of statistics" and cited paragraphs [0068]-[0069] and [0073] of Maltz for support. However, these cited passages are silent as to the claimed detection of a node which requires re-configuration. As already described by the Examiner, the teachings of Maltz describe collection of data, not detecting that a node requires re-configuration.

The above-reproduced arguments (incorporated herein) were originally presented on page 10, lines 5-13 of the First Response. The Examiner's response to these arguments is found in the twenty-second enumerated paragraph on page 8 of the Second Office Action in which the Examiner asserted the following:

In response to point (3), according to page 13, lines 10-18 of applicant's specification states "... the cluster of nodes can be monitored to identify when a change in configuration is appropriate. Changes in configuration can be appropriate when performance objectives are not met, when a node fails, when a node has become idle, when the workload has significantly changed, when a certain amount of time has passed, or upon the occurrence of any other suitable criteria" (i.e., detection of a node which requires re-configuration). In light of the specification, a node that requires reconfiguration is identified upon the occurrence of any suitable criteria. Similarly, Maltz teaches detecting a network element (node) which need scheduled collection/transmission of statistics (i.e., part of the reconfiguration process) upon an occurrence of criteria (e.g., certain amount of time has passed (i.e., schedule) or measured traffic is less than the mean traffic level (i.e., upon an occurrence of criteria)) ([0033], [0068]-[0069] and [0073]).

Appellants respectfully submit that the Examiner's characterization of the scope and content of the applied prior art is based upon a factually-unsupported inherency argument. The Examiner's analysis involves referring to certain portions of Appellants' teachings, which state that certain metrics may be monitored to determine if a node requires reconfiguration. The Examiner then reasons that if the applied prior art also teaches monitoring these same metrics,

then the applied prior art inherently teaches the limitation at issue (i.e., "detecting a node in the cluster which requires re-configuration.").

To be clear, paragraph [0033] does not teach the detection of a node which requires reconfiguration. Instead, paragraph [0033] of Maltz teaches collecting information, predicting bandwidth needs, and sending "control information back to the network elements to reconfigure the network 210" (emphasis added).

With regard to the Examiner's analysis that monitoring certain metrics, as taught by Maltz, inherently teaches detecting a node which requires reconfiguration, Appellants respectfully submit that such an inherency analysis is misplaced. Inherency may not be established by probabilities or possibilities. The mere fact that a certain thing may result from a given set of circumstances is not sufficient to establish inherency.³ To establish inherency, the extrinsic evidence must make clear that the missing function must necessarily be present in the thing described in the reference, and that the necessity of the feature's presence would be so recognized by persons of ordinary skill.⁴ This burden has not been met.

Page 13, lines 10-18 of Appellants' specification, which was cited by the Examiner, states the following:

FIG. 2 is a flow chart illustrating a process for autonomically optimizing individual nodes in a cluster. Beginning in block 205, the cluster of nodes can be monitored to identify when a change in configuration is appropriate. Changes in configuration can be appropriate when

³ In re Rijckaert, 9 F.3d 1531, 1534, 28 USPQ2d 1955, 1957 (Fed. Cir. 1993) (reversed rejection because inherency was based on what would result due to optimization of conditions, not what was necessarily present in the prior art); In re Oelrich, 666 F.2d 578, 581-82, 212 USPQ 323, 326 (CCPA 1981).

⁴ Finnegan Corp. v. ITC, 180 F.3d 1354, 51 USPQ2d 1001 (Fed. Cir. 1999); In re Robertson, 169 F.3d 743, 745, 49 USPQ2d 1949, 1950-51 (Fed. Cir. 1999); Continental Can Co. USA v. Monsanto Co., 20 USPQ 2d 1746 (Fed. Cir. 1991); Ex parte Levy, 17 USPQ2d 1461 (BPAI 1990).

performance objectives are not met, when a node fails, when a node has become idle, when the workload has significantly changed, when a certain amount of time has passed, or upon the occurrence of any other suitable criteria. In decision block 210, if a change in configuration is not appropriate, the process can return to block 205 and the cluster can continue to be monitored. Otherwise, the process can continue through block 215. (emphasis added)

As described in Appellants' specification, a change in configuration can be appropriate when certain criteria are met. However, this teach does not support a finding that a change in configuration must be made in response to the certain criteria being met. Since the Examiner has failed to set forth an factual evidence supporting a finding that the limitations at issue are explicitly taught by Maltz and since the Examiner has not established that this limitation is inherently (i.e., necessarily) disclosed by Maltz, Appellants maintain the argument that the Examiner has mischaracterized the scope and content of the teachings of Maltz.

The Examiner further asserted the following on page 5 of the First Office Action regarding claims 6 and 12:

identifying a workload hosted by said node ([0068]) (generating statistical summaries based on collected traffic information and storing the statistic in repositories (i.e. database)) and retrieving a set of configuration parameters associated with said workload ([0070], [0076], [0033]) (retrieved the data stored in repositories as input).

Statistical summaries of collected traffic information are not comparable to the claimed "workload." Moreover, Maltz does not teaching "retrieving a set of configuration parameters associated with said workload." Instead, as described in paragraph [0033], Maltz teaches computing network element configurations based upon (i) inputs that represent the traffic demand on the network, (ii) knowledge of network topology, and (iii) policy information. Computing does not identically disclose retrieving.

The above-reproduced arguments (incorporated herein) were originally presented on page 10, line 15 through page 11, line 2 of the First Response. The Examiner's response to these arguments is found in the twenty-third enumerated paragraph on pages 8 and 9 of the Second Office Action in which the Examiner asserted the following:

In response to point (4), on page 10, line 23 to page 11, line 2 of the remarks filed on 2/20/08, applicant states: "Moreover, Maltz does not teaching "retrieving a set of configuration parameters associated with said workload." Instead, as described in paragraph [0033], Maltz teaches computing network element configurations based upon (i) inputs that represent the traffic demand on the network, (ii) knowledge of network topology, and (iii) policy information. Computing does not identically disclose retrieving." Maltz teaches data (configuration parameters) stored in the TMS Statistic Repository 610 is used as an input to the TMS algorithm ([0076] and [0033]). Maltz further teach the data stored is associated with traffic information collected from the network element (i.e., workload) ([0068]). For example, traffic information is collected by measuring the number of bytes that flow out a line card interface. This means the data (configuration parameters) associated with an element at a host must be retrieved from repository in order to be used as input.

Although not explicitly stated, the Examiner's analysis begins by apparently construing the claimed "set of configuration parameters" as being comparable to "data stored in the TMS Statistic Repository 610." The Examiner, however, has failed to set forth any explanation that supports such a claim construction. The claimed set of configuration parameters are parameters that are directly used for configuring the node. However, the data stored in the TMS would not be considered by one having ordinary skill in the art as configuration parameters. Appellants recognize that although traffic data could be used to generate configuration parameters, traffic data itself are not parameters used to directly configure a node.

Referring to paragraph [0071], Maltz teaches that the TMS Statistics Repository 610 collects and analyzes ("processes, filters, compresses, and/or aggregates") traffic information. Referring to paragraph [0076], the data in the TMS Statistics Repository 610 can then be analyzed by the TMS Algorithm 200. Referring to paragraph [0033] and Fig. 2, Maltz teaches

1 computing network element configurations based upon (i) inputs that represent the traffic
2 demand on the network, (ii) knowledge of network topology, and (iii) policy information, and
3 computing control information does not identically disclose retrieving a set of configuration
4 parameters, which is what Appellants argued in the First Response.

5
6 Therefore, for the reasons stated above, Appellants maintain the argument that the
7 Examiner has again mischaracterized the scope and content of the teachings of Maltz.

8
9
10 The Examiner additionally asserted the following on page 5 of the First Office Action
11 regarding claims 6 and 12:

12 producing a new generation of configuration parameters based upon said retrieved set
13 using a computing process ([0033] and [0125]) (create configuration based on retrieved data
14 stored in repositories using an algorithm)
15

16 The Examiner yet again mischaracterizes the teachings of Maltz. Moreover, the Examiner's
17 analysis includes a logical inconsistency. As claimed, new configuration parameters are based
18 upon a retrieved set of configuration parameters. However, Maltz does not teach producing a
19 new generation of configuration parameters based upon the retrieved set of configuration
20 parameters. Instead, Maltz teaches producing network element configurations based upon (i)
21 inputs that represent the traffic demand on the network, (ii) knowledge of network topology, and
22 (iii) policy information. The (i) inputs that represent the traffic demand on the network, (ii)
23 knowledge of network topology, and (iii) policy information described by Maltz (i.e., the alleged
24 retrieved set) would not be considered by one having ordinary skill in the art as "old" versions of
25 "network element configurations." Thus, Maltz further fails to teach the limitations for which
26 the Examiner is relying upon Maltz to teach.

The above-reproduced arguments (incorporated herein) were originally presented on page 11, lines 4-20 of the First Response. The Examiner's response to these arguments is found in the twenty-fourth enumerated paragraph on page 9 of the Second Office Action in which the Examiner asserted the following:

In response to point (5), Maltz teaches configurations created for the network element (producing a new generation of configuration parameters) are based on output of the TMS algorithm ([0125]). As explained in point (2) above, the output of the TMS algorithm is based on the retrieved data from the repository (i.e., retrieved configuration parameter) as input to TMS algorithm ([0033] and [0068]). This means Maltz teaches producing a new generation of configuration parameters (creating new configuration based upon the output of TMS algorithm) based upon the retrieved set of configuration parameters (based on retrieved data from repository as input to TMS algorithm (i.e., old configuration)).

For ease of reference, paragraph [0033] of Maltz is reproduced below:

Turning again to the drawings, FIG. 2 is a block diagram of one presently preferred embodiment of the traffic management system (TMS). In this preferred embodiment, the traffic management system comprises a TMS Algorithm 200. The TMS Algorithm 200, which can be implemented with hardware and/or software, receives inputs that represent the traffic demand on the network 210. With these inputs and with knowledge of network topology and policy information, the TMS Algorithm 200 outputs network element configurations to automatically direct data based on the traffic demand. For example, the TMS can collect traffic information from all edge routers and switches in the network 210, predict bandwidth needs throughout the network 210, and send control information back to the network elements to reconfigure the network 210 to alter the forwarding of data so that network resources are better utilized (i.e., optimally utilized) based on the traffic demand on the network 210.

The flaws in the Examiner's analysis can easily be identified by comparing the input into the TMS Algorithm 200 (i.e., allegedly corresponding to the claimed retrieved set of configuration parameters) and the output from the TMS Algorithm 200 (i.e., allegedly correspond to the claimed new generation of configuration parameters). As already discussed above, an input into the TMS algorithm 200 is data in the TMS Statistics Repository 610, which is collected/analyzed traffic information. Referring to Fig. 2 of Maltz, other inputs include Network Topology Information, Network Policy Information, Explicit Allocation Request, and Service Level Agreements. On the other hand, the output from the TMS Algorithm 200 is

network element configurations. Thus, the data inputted into and the data outputted from the TMS Algorithm 200 are not data of the same type or character. Therefore, for the reasons stated above, Appellants maintain the argument that the Examiner has again mischaracterized the scope and content of the teachings of Maltz.

Regarding the Examiner's obviousness analysis, the Examiner asserted the following in the paragraph spanning pages 5 and 6 of the First Office Action:

It would have been obvious to one having ordinary skill in the art at the time of the invention was made to combine the teachings of Maltz and Nozawa because Nozawa's teaching of genetic computing process would enhance the configuration of Maltz's system by allowing genetic computing process to search for optimal configuration for the system.

The Examiner's analysis, however, ignores the teachings of Maltz. Specifically, the Examiner has failed to establish that Maltz teaches saving old configurations, which can then be used to create new configurations. Instead, Maltz teaches in paragraph [0033] that "the TMS Algorithm 200 outputs network element configurations to automatically direct data based on the traffic demand." Thus, the algorithm employed by Maltz is based upon "traffic demand," which is contemporaneous. On the contrary, a genetic computer process uses, in part, old information to create the new generation of configuration parameters. Maltz calculates the network element configurations based upon contemporaneous information, yet the Examiner's analysis has failed to explain why one having ordinary skill in the art would have been realistically impelled to modify Maltz so as to calculate the network element configurations based upon old configurations. Instead, the Examiner's analysis is conclusory without any factual basis.

The above-reproduced arguments (incorporated herein) were originally presented on page 11, line 22 through page 12, line 11 of the First Response. The Examiner's response to these arguments is found in the twenty-fifth enumerated paragraph on pages 9 and 10 of the Second Office Action in which the Examiner asserted the following:

In response to point (6), as explained in point (3) above, Maltz teaches creating configuration parameters (new configurations) based on input of data retrieved from repository (old configurations). Similarly, Nozawa's teaching of genetic computer process is also based in part of old information to create new generation. In addition, Maltz teaches other algorithm such as "genetic algorithm" can be used in the system ([0049]). Therefore, one having ordinary skill in the art would have been obvious to modify and to combine the teaching of Nozawa with Maltz.

The Examiner's analysis is predicated on the belief that Maltz teaches creating new configuration parameters based upon old configuration parameters. As already noted above, this belief is based upon an incorrect characterization of the teachings of Maltz. Thus, the Examiner's analysis, which flows from this belief, is flawed.

Conclusion

Based upon the foregoing, Appellants respectfully submit that the Examiner's rejections under 35 U.S.C. §§ 103, 112 based upon the applied prior art is not viable. Appellants, therefore, respectfully solicit the Honorable Board to reverse the Examiner's rejections under 35 U.S.C. §§ 103, 112.

To the extent necessary, a petition for an extension of time under 37 C.F.R. § 1.136 is hereby made. Please charge any shortage in fees due under 37 C.F.R. §§ 1.17, 41.20, and in connection with the filing of this paper, including extension of time fees, to Deposit Account 09-0461, and please credit any excess fees to such deposit account.

Date: November 12, 2008

Respectfully submitted,

/Scott D. Paul/

Scott D. Paul

Registration No. 42,984

Steven M. Greenberg

Registration No. 44,725

Phone: (561) 922-3845

CUSTOMER NUMBER 46320

VIII. CLAIMS APPENDIX

6. A method for autonomically optimizing a cluster of nodes, the method comprising the steps of:

detecting a node in the cluster which requires re-configuration;

identifying a workload hosted by said node and retrieving a set of configuration parameters associated with said workload;

producing a new generation of configuration parameters based upon said retrieved set using a genetic computing process; and,

reconfiguring said node with selected ones of said new generation of configuration parameters.

7. The method of claim 6, wherein said detecting step comprises the step of detecting at least one condition selected from the group consisting of a node crash, node idleness, node underperformance, and a change in workload hosted in said node.

8. The method of claim 6, wherein said producing step comprises the steps of:

performing a crossover operation for said configuration parameters in said retrieved set; and,

mutating at least one element of said configuration parameters in said retrieved set to produce the new generation of configuration parameters.

9. The method of claim 8, wherein said reconfiguring step comprises the steps of:

randomly selecting a new configuration from among said new generation of configuration parameters;

determining whether said randomly selected new configuration is viable; and,

reconfiguring said node with said randomly selected new configuration only if said new configuration is determined to be viable.

10. The method of claim 9, further comprising the step of writing said randomly selected new configuration to a knowledge base if said randomly selected new configuration is determined to be viable.

11. The method of claim 9, further comprising the steps of:

measuring node performance for said reconfigured node; and,

if said reconfigured node fails to meet baseline objectives for performance for said reconfigured node, selecting the new configuration for said node and performing said determining and reconfiguring steps for said selected new configuration.

12. A machine readable storage having stored thereon a computer program for autonomically optimizing a cluster of nodes, the computer program comprising a routine set of instructions for causing the machine to perform the steps of:

detecting a node in the cluster which requires re-configuration;

identifying a workload hosted by said node and retrieving a set of configuration parameters associated with said workload;

producing a new generation of configuration parameters based upon said retrieved set using a genetic computing process; and,

reconfiguring said node with selected ones of said new generation of configuration parameters.

13. The machine readable storage of claim 12, wherein said detecting step comprises the step of detecting at least one condition selected from the group consisting of a node crash, node idleness, node underperformance, and a change in workload hosted in said node.

14. The machine readable storage of claim 12, wherein said producing step comprises the steps of:

performing a crossover operation for said configuration parameters in said retrieved set; and,

mutating at least one element of said configuration parameters in said retrieved set to produce the new generation of configuration parameters.

15. The machine readable storage of claim 14, wherein said reconfiguring step comprises the steps of:

randomly selecting a new configuration from among said new generation of configuration parameters;

determining whether said randomly selected new configuration is viable; and,

reconfiguring said node with said randomly selected new configuration only if said new configuration is determined to be viable.

16. The machine readable storage of claim 15, further comprising the step of writing said randomly selected new configuration to a knowledge base if said randomly selected new configuration is determined to be viable.

17. The machine readable storage of claim 15, further comprising the steps of:
measuring node performance for said reconfigured node; and,
if said reconfigured node fails to meet baseline objectives for performance for said reconfigured node, selecting the new configuration for said node and performing said determining and reconfiguring steps for said selected new configuration.

IX. EVIDENCE APPENDIX

No evidence submitted pursuant to 37 C.F.R. §§ 1.130, 1.131, or 1.132 of this title or of any other evidence entered by the Examiner has been relied upon by Appellants in this Appeal, and thus no evidence is attached hereto.

X. RELATED PROCEEDINGS APPENDIX

Since Appellants are unaware of any related appeals and interferences, no decision rendered by a court or the Board is attached hereto.